

The Global Lambda Visualization Facility: Technologies for Ultra-High-Definition Wide-Area Visualization and Collaboration

Jason Leigh, Luc Renambot, Andrew Johnson, Byung-il (Brent) Jeong, Ratko Jagodic, Nicholas Schwarz,
Dmitry Svistula, Rajvikram Singh, Venkatram Vishwanath, Brenda Lopez, Dan Sandin, Tom Peterka,
Javier Girado, Robert Kooima, Jinghua Ge, Lance Long, Alan Verlo
Electronic Visualization Laboratory, University of Illinois at Chicago

Donna Cox, Robert Patterson, David Semeraro, Stuart Levy
National Center for Supercomputing Applications / University of Illinois at Urbana Champaign

Jonas Talandis, Joe Reitzer, Tom Prudhomme
Technology, Research, Education, and Commercialization Center / NCSA

Tom Coffin
ACCESS Center / NCSA

Brian Davis
Earth Resources Observation Systems (EROS), US Geological Survey (USGS)

Paul Wielinga, Bram Stolk
SARA Computing and Networking Services

Jaeyoun Kim, Sangwoo Han, JongWon Kim
Gwang-ju Institute of Science and Technology

Brian Corrie, Todd Zimmerman
Simon Fraser University

Pierre Boulanger
University of Alberta

1 Abstract

The research outlined in this paper marks an initial global cooperative effort between visualization and collaboration researchers to build a persistent virtual visualization facility linked by ultra-high-speed optical networks. The goal is to enable the joint research and development of the necessary hardware, software and interaction techniques to realize the next generation of end-user tools for scientists to collaborate on the global Lambda Grid. This paper outlines some of the visualization research projects that were demonstrated at the iGrid 2005 workshop in San Diego, California.

2 Introduction

In a decade's time, high-performance computing has proven its value in science, medicine, engineering, and even the arts. These data-intensive domains rely on high-performance computing to process the terabytes of raw data (for example, from sensors) to produce meaningful insight, often enabled by high-quality visualizations. The need for the computing infrastructure to support collaborative work between distributed users has also grown, as research has increasingly become multi-disciplinary endeavors.

Up until the last five years however, the predominant model for supporting data-intensive collaborative science and engineering involved replicating and co-locating the computing resources with the storage systems, and then transmitting the highly distilled results of the computation to expensive remote visualization systems at each collaborator's site. This was because bandwidth was a scarce commodity and so it was more efficient to co-locate the data with the computation than to move the data, on-demand, to the computation. This meant that computer centers not only had to maintain clusters of computers for number-crunching, they also had to maintain backed-up storage facilities to keep copies of the raw data for their users. Recently, with rate of decline of the cost of bandwidth far exceeding that of computing and storage, it has become viable for scientists to connect to ultra-high-speed networks, more cheaply than they can build and maintain large computing, storage, and visualization systems¹[Stix01]. This is the fundamental premise behind the concept of shared cyber-infrastructure, and is being deeply researched amongst the international partners in the OptIPuter project[Smarr03]. By bridging high-performance computing systems with ultra-high-speed networks that can approach the capacity of computer system buses, one can potentially reduce the need to duplicate systems except where redundancy (for example, for backup) is important. In the same way that users today search for information on the Web with a browser and rely on Google to maintain the vast clusters of computers that organize information on the Internet, scientists who could not ordinarily afford high-performance computing, can collaborate and access the network of cyber-infrastructure (the Lambda Grid) using low-cost, scalable, high-resolution network-aware displays and high-definition video conferencing[Renambot05].

The research outlined in this paper and demonstrated at the iGrid 2005 workshop² at CALIT2 in San Diego, California, marks an initial global cooperative effort between visualization and collaboration researchers to build a persistent virtual distributed facility (GLVF- the Global Lambda Visualization Facility) to enable the joint research and development of the hardware, software and interaction techniques, that are to become the primary interfaces between domain scientists and the Lambda Grid. The institutions currently participating in GLVF include:

- The Electronic Visualization Laboratory at the University of Illinois at Chicago;
- The California Institute for Telecommunications and Information Technology (CALIT2);
- The National Center for Supercomputing Applications (NCSA) at the University of Illinois Urbana-Champaign, and the ACCESS Center in Washington DC;
- Earth Resources Observation Systems (EROS), US Geological Survey (USGS);
- The Technology, Research, Education, and Commercialization Center in Dupage County, Illinois;
- SARA Computing and Networking Services in Amsterdam, the Netherlands;
- Gwang-ju Institute of Science and Technology (GIST) in Gwang-ju, South Korea;
- Korean Institute of Science and Technology Information (KISTI) in Daejeon, South Korea;
- Simon Fraser University and the University of Alberta in Canada.

¹ Gary Stix, The Triumph of the Light, Scientific American, January 2001- bandwidth is doubling every 9 months; computing at every 18 months; and storage at every 12 months.

² See editor's article in this issue of the Future Generation Computer Systems Journal for a description of iGrid 2005.

GLVF sites are all equipped with a wide-range of visualization technologies including high-definition passive stereo displays, fully immersive CAVEs, auto-stereoscopic displays, and ultra-high resolution tiled projection and LCD displays, and high-definition cameras. In most cases these systems are already connected to high-speed multi-gigabit metropolitan, national and international networks. Internationally, GLVF leverages the networking infrastructure that has been amassed as part of the Global Lambda Integrated Facility (GLIF- www.glif.is). This enables GLVF to become a one-of-a-kind resource for conducting long-term systematic research in network- and data-intensive visualization and collaboration. For domain scientists, GLVF can provide a means for them to learn how to apply Lambda Grid technology to transform their own areas of research.

3 GLVF at iGrid

At the iGrid workshop GLVF researchers streamed real-time ultra-high-resolution graphics, and high-definition video from their home sites for display on a variety of display systems including: a 100-Megapixel LambdaVision tiled LCD display (www.evl.uic.edu/cavern/lambdavis); a passive stereo XGA-resolution GeoWall display (www.geowall.org); a passive stereo high-definition display; and the autostereoscopic Varrier display [Sandin05].

3.1 The Scalable Adaptive Graphics Environment

Luc Renambot (luc@mac.com), Byung-il Jeong, Ratko Jagodic, Rajvikram Singh, Andrew Johnson, Jason Leigh - Electronic Visualization Laboratory (EVL), University of Illinois at Chicago

The OptIPuter project envisions that in the future it will become routine for users to work and collaborate in rooms whose walls are made from seamless ultra-high-resolution displays that are fed by data streamed over ultra-high-speed networks from distantly located visualization, storage servers, and high-definition video cameras [Smarr03, Leigh03]. EVL has taken the first steps towards this vision by building LambdaVision – a 17-foot wide, tiled display built from an 11x5 array of LCD screens with a total resolution of 100 megapixels. LambdaVision is primarily enabled by the Scalable Adaptive Graphics Environment (SAGE) [Renambot04] a middleware system for managing visualization and high-definition video streams that are presented on the ultra-high-resolution display. Each visualization application (such as 3D rendering, remote desktop, video streams, large 2D maps) streams its rendered pixels (or graphics primitives) to SAGE, allowing for any given layout onto the displays (e.g. the output of arbitrary M by N pixel rendering cluster nodes can be streamed to X by Y pixel display screens).

The graphics streaming architecture of SAGE addresses two non-trivial problems in scientific visualization. One is heterogeneity: since most visualization applications are closely tied to their graphics environment, it is difficult to integrate various visualization applications into a unified graphics environment. The other is scalability: the ability of visualization software and systems to scale in terms of the amount of data they can visualize and the resolution of the desired visualization [Blanke00]. SAGE addresses the heterogeneity problem by decoupling graphics rendering from graphics display so that visualization applications developed on various environments can easily migrate into SAGE by streaming their pixels into the virtual frame buffer. Also, SAGE provides scalability by supporting any number of rendering and displaying nodes, number of tiles, and screen resolution, and the SAGE visualization applications have extremely fast access to huge databases at remote or local sites taking advantage of affordable ultra-high-bandwidth networks.

At the iGrid 2005 workshop, EVL succeeded in simultaneously streaming a wide variety of content including: real-time parallel volume rendered images of the Visible Korean Human from KISTI (over KREONet2); NCSA's pre-rendered high-definition movies of tornado simulations from SARA (over SURFNet); and 1-foot resolution aerial photographs of New Orleans in the aftermath of the Katrina hurricane, and live high-definition video from EVL (over CAVEWave/National Lambda Rail). 10Gb/s was allocated for the experiment of which 9Gb/s was used.

3.2 TOPS - Technology for Optical Pixel Streaming

Paul Wielinga (wielinga@sara.nl), Bram Stolk – SARA Computing and Networking Services

SARA has been studying the separation of the (virtual) graphics card and the display in the context of the OptIPuter principle. Our challenge is to cope with latency on long-distance connections. SARA's approach to driving tiled displays from remote render clusters is a lightweight approach. Our key insight into this problem is based on the short lifespan of a picture element in a video stream. At 50 frames per second, a picture element, or pixel, is only visible for 20 ms. If such a pixel is lost during communication, chances are it will remain unnoticed, because apart from spatial coherence, video streams also exhibit temporal coherence. Our approach therefore uses unreliable network transport (UDP) as opposed to a reliable (TCP) transport. This will give a valuable benefit, as high throughput is easily realized when acknowledgements from remote sites with large latencies are no longer required.

The viability of our approach was successfully demonstrated during iGrid2005, where ultra-high-resolution images were streamed live from Amsterdam to San Diego over two 10Gb/s connections, whereby the application consumed 18 Gb/s bandwidth, with peaks of 19.5 Gb/s. The application was controlled from the display side, and the low bandwidth control signals were sent reliably over TCP to Amsterdam. The high-bandwidth stream, representing ultra-high-resolution imagery, was displayed on the LambdaVision.

Bandwidth usage was controlled by the observer, by varying frame rate on the render cluster. Initially, the frame rate was limited by the receiving side, as the large amount of small 1500 byte packages locked up the CPU's of the receiving display nodes. The total connection between the render cluster and the display cluster was later jumbo-frame enabled. When using an MTU size of 9000 bytes, we could fill all available bandwidth. The frame rate on the panel was increased to 9 frames/second, creating a data stream of 2,3 GBytes/second. This is the equivalent of transferring a complete DVD in 2 seconds over 5000 miles distance.

3.3 Improved High-Quality Media Service Architecture for Multi-Party Collaboration Environments

Jaeyoun Kim, Sangwoo Han, and JongWon Kim ([jykim,swhan,jongwon}@nm.gist.ac.kr](mailto:{jykim,swhan,jongwon}@nm.gist.ac.kr)) - Networked Media Lab (NML), Dept. of Information and Communications, Gwangju Institute of Science and Technology (GIST, formerly known as K-JIST)

With the advent of Gb/s-range Research and Engineering networks, it is now possible to accommodate bandwidth-demanding high-quality (DV / HDV / Stereo HD / Uncompressed HD) video services into computer-mediated collaboration environments such as the AccessGrid (AG)[AG]. In such environments, the enhanced media service is expected to improve Quality of Experience (QoE) of participants in a collaboration session. Thus, NML designed and implemented an architecture for high-quality video services for multi-party collaborative environments, such as the AG. This attempt to enable high-quality media services for the AG began in 2003. In 2004, the development of DV/HDV support for the AG was released for public use (<http://ace.nm.gist.ac.kr>). NML designed extensions for the AGTk (AG toolkit) video services, so called "ExtendedVideo{Producer,Consumer}Services" by considering such issues as flexible multicast address allocation, versatile video codec/application support, and negotiation of service-capability. However, even though the AGTk augmentation enabled DV/HDV support, two key missing capabilities prompted NML to develop a new architecture, named AG Media / SVC (Scalable Visualization Consumer) [Han05, Kim05]. The first capability that the new architecture addresses is the automatic negotiation of common streaming media formats (DV / HDV / Stereo HD / Uncompressed HD) between remote sites that are participating in a collaboration. The second much needed capability is a means to decouple CPU-intensive video decoding from image display so that the video can be easily manipulated when presented on displays of arbitrary resolution. During the summer of 2005, NML developed a version of SVC that used EVL's SAGE to distribute high-definition video to EVL's LambdaVision display. This

was successfully demonstrated at iGrid 2005 by streaming video from Chicago to San Diego over CAVEWave.

3.4 BPlayer – High-definition Stereoscopic Graphics Streaming

Donna Cox (cox@ncsa.uiuc.edu), Robert Patterson, Patrick Dorn, Paul Wefel, Stuart Levy – National Center for Supercomputing Applications

Research projects at the University of Illinois in Urbana and Chicago have been brought together to accomplish the goal of streaming high-definition stereo movies long distances across optical networks. At iGrid 2005, visualizations of several scientific simulations were streamed from the National Center for Supercomputing Research in Urbana Illinois to the iGrid facility in San Diego California. Stereoscopic 3D, time-dependent, data were rendered to individual images at high-definition resolution. These images were then played back in sequence and streamed across a high-speed optical network to a stereoscopic HD projection system at iGrid. NCSA's BPlayer image playback software was augmented with EVL's SAGE to allow it to stream the playback over high-speed networks.

BPlayer is an application for displaying movies stored on disk as sequences of images. BPlayer software is used at NCSA to display HD content on a passive stereo display system driven by a single dual processor computer with a high performance hardware raid. Performance of this movie playing techniques depends on the ability to read image data from disk at a high rate. BPlayer uses a multi-threaded staggered sequential disk access scheme in order to realize the full bandwidth of the disk array. The integration of SAGE with BPlayer enabled it to stream images across the network to a remote display system. This combination of software was able to support the display of HD movies at various resolutions. The iGrid demonstration achieved approximately 14 frames per second display rate and used approximately 1.4 Gb/s. While not entirely satisfactory the frame rate was adequate given the distances involved. Further work is necessary to identify bottlenecks and address them.

3.5 Auto-Stereoscopic Video Conferencing

Dan Sandin (dan@uic.edu), Tom Peterka, Javier Girado, Robert Kooima, Jinghua Ge – Electronic Visualization Laboratory, University of Illinois at Chicago

Prior teleconferencing experiences have largely been a 2D experience, but this changed at iGrid 2005, as conference attendees sat and discussed the benefits of *autostereo* teleconferencing “face-to-face” in real-time 3D with scientists at EVL. This essentially meant that viewers could see stereoscopic 3D video of a remote participant without the need to wear special glasses. The display system consisted of an EVL-designed device called the Varrier auto-stereo display [Sandin05] driven by a small cluster of computers to track the viewer and render a proper view. Initial feedback from iGrid was positive and participants certainly appeared to enjoy the interaction. The CAVEwave 10 gigabit optical networking between UCSD and UIC facilitated audio and stereo video streaming in real-time, using approximately 220 Mb/s of bandwidth in each direction. End to end latency was approximately .3 s (one way), and autostereo images of participants were rendered at 25 Hz. Raw, unformatted YUV411 encoded images were transmitted, using one half the bandwidth as compared to a corresponding RGB format. The real-time display for the most part appeared flawless even though approximately 10% of packets were dropped. The robustness is due to the high transmission and rendering rates. The sending rate was 30 Hz, while the rendering rate was approximately 25 Hz, limited mainly by pixel fill rates and CPU loading. The net result is that the occasional lost packets were negligible. The excellent network quality of service can be attributed to the limited number of hops and routers between source and destination. There were exactly 4 hosts along the route, including source and destination; so only two intermediate routers were encountered. The success of the teleconferencing is directly attributable to the 10 Gb/s CAVEwave light-path over the National LambdaRail, as the application simply sends raw stereo video and audio data and relies on the network to successfully deliver quality of service.

3.6 The Solutions Server

Brian Corrie (bcorrie@sfu.ca), Todd Zimmerman – Simon Fraser University; Pierre Boulanger – University of Alberta

The Media LightPath (MLP) project is a CANARIE funded project that is exploring the use of user-controlled lightpaths (UCLP) [Wu03] for advanced collaboration and visualization. One of the goals of the MLP project is to provide a UCLP-based collaboration and visualization capability to collaborating institutions that have lightpath capability. As a first step towards this end, the MLP project participated in the iGrid 2005 GLVF demonstration. Utilizing the Solutions Server, a package for coupling computational simulation with remote visualizations developed at the University of Alberta, we ran a computational simulation at the University of Alberta in Edmonton. The simulation streamed computational results over the WestGrid core network to a visualization resource at Simon Fraser University in Vancouver. The visualization machine then rendered this computational data as a stereo pair of images, sent these images to San Diego, and displayed them on a dual-projector passive stereo display at the iGrid venue. This data was sent over a dedicated gigabit lightpath from Vancouver to San Diego using the VizServer package developed by SGI.

The main result that MLP researchers took away from their iGrid experiences was the large impact that distance can have on software performance. Before iGrid, MLP had conducted some limited testing on a temporary lightpath between Vancouver and Edmonton as well as using the shared WestGrid lightpath. Throughput over the local campus backbone exceeded 600 Mbps. Experiments over a test lightpath and the WestGrid network decreased application-level performance to approximately 200 Mbps. When the same experiment was extended to iGrid in San Diego, that performance further decreased to a mere 20 Mbps. It was believed that the fundamental problem lay in the inadequate tuning of the end systems. Furthermore it was clear that the tuning had to be done on a case-by-case basis and clearly thorough research is needed to develop an agile solution that could account for all networking conditions. This could be achieved if the networking and visualization infrastructure remained persistent well beyond the time span of the workshop.

4 Future Challenges

iGrid 2005 proved to be a successful launching point for the Global Lambda Visualization Facility. The collaborating researchers found it highly informative to be able to evaluate their own techniques and to compare them against alternative approaches developed by their colleagues. The next goal is to ensure the permanence of the testbed so that the level of cooperation that was initiated by iGrid can be further expanded. An immediate challenge that the group will collectively solve is a means to provide high-definition video and ultra-high-resolution visualization multicasting between the sites. Multicasting at the magnitude of tens of gigabits is challenging both in terms of networking and the end-systems that receive the multicast streams. Other challenges include the development of a mechanism for sites to reserve visualization resources for conducting wide-area visualization experiments. We anticipate the need to leverage much of the existing Grid middleware technology to support this.



Figure 1: Top Left: SAGE showing several high resolution streams (high-definition thunderstorm simulation; high-definition live video; ultra-high-resolution aerial photography of New Orleans after the Katrina Hurricane) simultaneously on the LambdaVision display; Bottom Left: TOPS showing a single high resolution montage of Amsterdam; Top Right: Autostereoscopic video conferencing on the Varrier display; Middle Right: Solutions Server streaming graphics to a passive stereo XGA resolution display; Bottom Right: BPlayer streaming high-definition stereo movies to a high-definition passive stereo display.

5 Acknowledgments

The Electronic Visualization Laboratory (EVL) at the University of Illinois at Chicago specializes in the design and development of high-resolution visualization and virtual-reality display systems, collaboration software for use on multi-gigabit networks, and advanced networking infrastructure. These projects are made possible by major funding from the National Science Foundation (NSF), awards CNS-0115809, CNS-0224306, CNS-0420477, SCI-9980480, SCI-0229642, SCI-9730202, SCI-0123399, ANI-0129527 and EAR-0218918, as well as the NSF Information Technology Research (ITR) cooperative agreement (SCI-0225642) to the University of California San Diego (UCSD) for "The OptIPuter" and the NSF Partnerships for Advanced Computational Infrastructure (PACI) cooperative agreement (SCI-9619019) to the National Computational Science Alliance. EVL also receives funding from the State of Illinois, General Motors Research, the Office of Naval Research on behalf of the Technology Research, Education, and Commercialization Center (TRECC), and Pacific Interface Inc. on behalf of NTT Optical Network Systems Laboratory in Japan. The GeoWall, GeoWall2, Personal GeoWall2 (PG2), and LambdaVision are trademarks of the Board of Trustees of the University of Illinois.

SARA's TOPS research is funded in part by the GigaPort NG project. SARA is the Netherlands' High Performance Computing and Networking Center. It operates the SURFnet NOC and houses SURFnet5, SURFnet6 and NetherLight cores and the Netherlands' national Supercomputers.

Simon Fraser University and the University of Alberta's MediaLight Path and Solution Servers projects are supported by CANARIE, SGI, CFI, and BCKDF.

KISTI (Korea Institute of Science and Technology Information), GIST (Gwangju Institute of Science and Technology) are supported through the BK-21 program, and EVL at UIC.

6 References

[AG] The AccessGrid – www.accessgrid.org

[Stix01] Stix, G., "The Triumph of the Light," Scientific American, January 2001.

[Renambot05] Renambot, L., Johnson, A., Leigh, J., "LambdaVision, Building a 100 Megapixel Display," NSF CISE/CNS Infrastructure Experience Workshop, Champaign, IL – July 27, 2005.

[Smarr03] Smarr, L., Chien, A. A., DeFanti, T., Leigh, J., Papadopoulos, P. M., "The OptIPuter" Communications of the ACM, Volume 46, Issue 11, November 2003, pp. 58-67.

[Leigh03] Leigh, J., Renambot, L., DeFanti, T. A., et al, "An Experimental OptIPuter Architecture for Data-Intensive Collaborative Visualization", 3rd Workshop on Advanced Collaborative Environments, Seattle, WA, June 2003

[Blanke00] Blanke, W., Bajaj, C., Fussell, D., Zhang, X., "The Metabuffer: a Scalable Multiresolution Multidisplay 3-D Graphics System using Commodity Rendering Engines." Tr2000-16, University of Texas at Austin, February 2000.

[Renambot04] Renambot, L., Rao, A., Singh, R., Jeong, B., Krishnaprasad, Naveen, Vishwanath, V., Chandrasekhar, V., Schwarz, N., Spale, A., Zhang, C., Goldman, G., Leigh, J., Johnson, A., "SAGE: the Scalable Adaptive Graphics Environment," Proceedings of WACE 2004, Nice, France, 09/23/2004 - 09/24/2004.

[Sandin05] Sandin, D., Margolis, T., Ge, J., Girado, J., Peterka, T., DeFanti, T., "The Varrier Autostereoscopic Virtual Reality Display," ACM Transactions on Graphics, Proceedings of ACM SIGGRAPH 2005, vol 24, no 3, pp. 894-903

[Wu03] Wu, J., Campbell, S., Savoie, J. M., Zhang, H., Bochmann, G. v., St.Arnaud, B., "User-managed end-to-end lightpath provisioning over CA*net 4," Proceedings of the National Fiber Optic Engineers Conference (NFOEC), Orlando, FL, USA, Sept 7-11, 2003, pp. 275-282.

[Han05] Han, S., Kim, JW., "Adaptive High-quality Video Service for Network-based Multi-party Collaboration", in Proc. of SPIE ITCOM'2005: Multimedia Systems and Applications VIII, Boston, Oct. 2005.

[Kim05] Kim J., Kim JW., "Decomposable Decoding and Display Structure for Scalable Media Visualization over Advanced Collaborative Environments", in Proc. of SPIE ITCOM'2005: Multimedia Systems and Applications VIII, Boston, Oct. 2005.